

AD-A250 536



ENTATION PAGE

Form Approved
OPM No. 0704-0188

average 1 hour per response, including the time for reviewing instructions, searching existing data sources gathering and Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions are for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to and Budget, Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE		3. REPORT TYPE AND DATES COVERED	
		1989		Unknown	
4. TITLE AND SUBTITLE				5. FUNDING NUMBERS	
Beyond Specificity				DAAB10-86-C-0567	
6. AUTHOR(S)					
Henry E. Kyburg, Jr.					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)				8. PERFORMING ORGANIZATION REPORT NUMBER	
University of Rochester Department of Philosophy Rochester, NY 14627					
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
U.S. Army CECOM Signals Warfare Directorate Vint Hill Farms Station Warrenton, VA 22186-5100				92-TRF-0007	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION/AVAILABILITY STATEMENT				12b. DISTRIBUTION CODE	
Statement A; Approved for public release; distribution unlimited.					
13. ABSTRACT (Maximum 200 words)					
<p>A number of writers have suggested that specificity can be called upon to adjudicate competing default inferences. In the foundations of statistics, specificity is one of several ways to adjudicate the claims of competing reference classes. This suggests that in default inferences also other principles than specificity may be needed. This paper gives examples substantiating this suggestion, and provides formulations of the few other principles needed.</p>					
14. SUBJECT TERMS				15. NUMBER OF PAGES	
Artificial Intelligence, Data Fusion, Specificity				8	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT		18. SECURITY CLASSIFICATION OF THIS PAGE		19. SECURITY CLASSIFICATION OF ABSTRACT	
UNCLASSIFIED		UNCLASSIFIED		UNCLASSIFIED	
				20. LIMITATION OF ABSTRACT	
				UL	

BEYOND SPECIFICITY*

(area: B6)

by
Henry E. Kyburg, Jr.
University of Rochester

kyburg@cs.rochester.edu

abstract

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

A number of writers have suggested that specificity can be called upon to adjudicate competing default inferences. In the foundations of statistics, specificity is one of several ways to adjudicate the claims of competing reference classes. This suggests that in default inferences also other principles than specificity may be needed. This paper gives examples substantiating this suggestion, and provides formulations of the few other principles needed.

1. It has been suggested (Poole, 1985; Touretzky, 1984, 1987; Neufeld, 1988; Bacchus, 1988; Etherington, 1987) that considerations suggested by probability theory may throw light on non-monotonic inference. The problems of non-monotonic inference and probability do seem to be very close to each other; it is the purpose of this paper to explore that relation further. It may seem that we are using probabilistic considerations to throw further obscurity on non-monotonicity. To avoid this impression, we shall first present intuitive cases of non-monotonic inference, without reference to probability, and only subsequently point out the connections.

The general nature of the problem we are considering is the following: We have a set of premises in our body of knowledge or knowledge base, from which we would ordinarily expect to be able to infer a certain statement *S*. But there is another set of statements, that may equally well be regarded as being part of our knowledge base, in that same situation, from which we could infer the denial of *S*. In many cases, what we suppose ourselves to know in the first place entails that these conclusion upsetting statements are part of our knowledge. (Note: this is not just a matter of not being able to infer *S*, but a matter of being able to infer the denial of *S*.)

The classical case is that of Tweety the penguin. We want to infer that Tweety does not fly, even though we know at the same time (ipso facto, we might even say!) that Tweety is a bird and that typically birds fly. By themselves, these facts would warrant the opposite conclusion, namely: Tweety flies.

One approach to this problem, suggested in various forms by the authors cited above, is to observe that when these two possible arguments clash, we prefer the argument with the "most specific" premises. In this case that specificity picks



92-13709
■■■■■■■■■■

out the argument that classifies Tweety as a penguin rather than the argument that merely classifies Tweety as a bird. This situation is represented by figure 1.

Historically it is worth noting that the idea of specificity goes back at least to Reichenbach (1949). The term "specificity" was employed by Hempel (Hempel, 1968), and there has been some philosophical discussion of the notion. It is also worth observing that exceptions -- even singleton exceptions -- are conveniently handled by reference to specificity. If Tweety is the only penguin in the world who can fly, and we happen to know it, then the facts that Tweety is in {Tweety} (or has the property of being Tweety), and that all the members of {Tweety} (or everything with the property of being Tweety) flie(s), lead us to the conclusion that Tweety flies after all.

We may present the problem more formally this way: We have a knowledge base containing premises

P_1, P_2, \dots, P_n

On this basis we want to obtain the conclusion C . But if our knowledge base contains the P_i , it also contains

R_1, R_2, \dots, R_k

either because they are implied by the P_i or because (like "Birds typically fly") they represent natural assumptions. But given the R_i in our knowledge base, in the absence of the P_i , we would conclude the denial of C , $\neg C$.

2. Specificity solves examples of the form illustrated in figure 1. But even minor variations call for something more. Even in the case of Tweety, this can be seen. If we know that Tweety is a penguin, and that penguins don't typically fly, then we also know

(1) "{Tweety} is a subset of birds."

And in general we know that

(2) "Typically subsets of birds are also subsets of flying objects."

From which it is natural to infer that Tweety flies.

Now this may be thought to be a strange and unnatural way of expressing our knowledge about the ability of birds to fly. No doubt. But it represents a fairly straight-forward logical translation of the R -premises of the first example. (If one didn't like the set of which Tweety is the only member, one could talk about the property of being Tweety.) True, our knowledge base cannot be closed under logical implication, but it seems artificial to rule out any particular forms of inference as illegitimate. How would we draw the line? If we are looking to allow non-monotonic inferences, we should surely allow some simple deductive inferences as well.

So it is not unreasonable to suppose that (1) and (2) are in our knowledge base. What prevents the inference to {Tweety} is a subset of the flyers?

In the original form of Tweety's problem, specificity did it. That won't work here, since penguins are not a subset of sets of birds. (Sets of anything are abstract objects; penguins aren't.) But we can employ almost the same principle here, as illustrated in figure 2. Here is a rough statement of the specificity principle that takes care of the two cases we have considered so far.

If "A 's are typically C 's" and "x is an A " are in our knowledge base, and "B 's are typically D 's" and "y is a B " are in our knowledge base, and "x is a C if and only if y is not a D " is in our knowledge base, then the first inference is to be preferred to the second whenever we also know that there is a subset B^* of B such that "y is in B^* " and "Typically members of B^* are not members of D " are in our knowledge base.

This rule applies to the first example, with penguin for $B^* = A$, bird for B , flies for C , and $x = y = \text{Tweety}$. It applies to the second example with non-fliers for D , and {Tweety} for y .

3. The "artificial" form of the Tweety example is rather baroque, and unlikely to arise except in the mind of a perverse logical or a perfectly logical computer program. The second counter-example to simple specificity, represented in figure 3, is much more natural and in can no way be construed as a matter of "specificity."

Here we know that a room contains ten cages, nine of which contain one healthy sparrow and two fat penguins, and the tenth cage containing 171 sparrows and only two penguins. We select a cage, and then a bird (whom we call Tweety) from the cage. Typically, this bird will not be a flier. But note: we also know that the selected bird is a bird in the room, and typically birds in that room do fly.

Observe that there is no subset (no specification) of the set of birds in the room to which we know that Tweety, the selected bird, belongs, and in which the typical bird is a non-flier. "For all we know," Tweety was selected from the tenth cage.

It is no answer to say that "most of the time" the selected bird will have come from one of the other cages, whatever "most of the time" means here. We can adjust the numbers of cages and numbers of each kind of bird in each cage so that it is not the case that "most of the time" we will select the cage with typical non-fliers, but it is the case that "most of the time" we will select birds that do not fly.¹

A general form on an appropriate rule is this:

If "A 's are typically C 's" and "x is an A " are in our knowledge base, and "B 's are typically D 's" and "y is a B " are in our knowledge base, and "x is a C if and only if y is not a D " is in our knowledge base, then the first inference is to be preferred to the second whenever we can find a cross product $B^* \times B$, a pair $\langle z, y \rangle$, and a predicate of pairs, D^* , of which the following are known to be true in our data base:
 $\langle z, y \rangle$ is D^* just in case y is D ; $\langle z, y \rangle$ is in $B^* \times B$; our knowledge about $B^* \times B$ and D^* matches our knowledge about B and D ; and, finally, there is a subset E of $B^* \times B$ such that $\langle z, y \rangle$ belongs to it,

and our knowledge about E and D^* matches our knowledge about A and C .

In the case at hand, C is the property of pairs consisting of a cage and a bird that holds when the bird can't fly. B^* is the set of cages. $E = A$ is the subset of $B^* \times B$ that satisfies the condition that the bird part of the pair comes from the cage part of the pair.

4. The final form of inference to be considered is a bit more specialized than the preceding two, but again cannot be explained in terms of specificity. The situation is illustrated in figure 4.

Suppose we have examined a sample of 10,000 birds, and observed that 75% of them are fliers. Under the right conditions, it is reasonable for us to conclude that more than 70% of birds fly, since under the right conditions samples of 10,000 typically represent the populations from which they are drawn.

But then we also have in our knowledge base knowledge of a sample of 5,000 birds, of which only 50% are fliers. Under the right conditions, it would be reasonable to conclude that less than 70% of birds fly, since samples of 5,000 typically represent the populations from which they are drawn.

As we have told the story, if the "right conditions" have been met for the first inference, we want to be able to show, that the "right conditions" cannot be met for the second inference. We want to prefer the first inference. Again, "specificity" doesn't help. But the fact that the sample of the second inference is included in the sample of the first inference provides a reasonable criterion.

We can state the rule as follows:

If " A 's are typically C 's" and " x is an A " are in our knowledge base, and " B 's are typically D 's" and " y is a B " are in our knowledge base, and " x is a C if and only if y is not a D " is in our knowledge base, then the first inference is to be preferred to the second if it is possible to construe the two inferences as statistical inferences differing only in that the sample of the second inference is a subset of the sample of the first inference.

The particular case under discussion corresponds to this rule in an obvious way.

5. Are there other forms of competition among plausible inferences than those mentioned here? I think not. My evidence is that the three forms of preferences just illustrated (actually, two, since "specificity" can be construed as a special case of the cross product construction) are the only forms that seem to be necessary in accounting for the choice of a reference class in an epistemic probability theory. See Kyburg (1983, 1988) for more details. There may be other structures that should be taken account of that have not yet been noticed. But for the moment, at least, these three correspond to our most clear-cut intuitions.

It may be observed that "typicality" and "frequency" are different ideas, and may not be subject to the same rules and constraints. This may be so, but there is some reason to think that they cannot conflict too severely.

In the first place it is clear that we cannot claim that A 's are typically C 's, when most A 's aren't C 's -- at least without a long and fairly complicated story. We might suppose that x is an A , that most A 's are C 's, and at the same time that x is a B , while B 's are not typically C 's. The upshot of this knowledge would depend on the relation between A and B . If it is one of those relations addressed above, then the considerations adumbrated there should determine the outcome rather than a contest between "typicality" and "frequency".

Another problem that may seem worrisome is whether the process of specification, Bayesianization, or sample expansion will always end. Specification and sample expansion clearly do: it is no big deal to constrain our references to finite sets; and no finite set can admit of arbitrarily fine specification. Before long we must come to that most specific class: that class of which the object in question is the only member. Analogously, our samples can only be so big. There is a largest, and we may assume that our knowledge base knows about it. The problematic case concerns cross product formation. But even this case cannot lead us on indefinitely; the complexity of any compound experiment must be bounded.

We thus conclude, tentatively, that the enumerated considerations are all the ones that are relevant to the adjudication of the claims of competing non-monotonic inferences. We also conclude, definitively, that specificity, even if construed quite broadly, doesn't do the trick. And we conclude, quite generally, that however "typicality" is construed, statistical considerations can throw light on inferences that depend on typicality.

notes.

* Research underlying this paper was supported in part by the U. S. Army Signals Warfare Center.


1. Suppose "most of the time" means with a relative frequency of greater than $1 - 10^{-n}$ of the time. Then we need 10^n cages, in all but one of which we have 10^{n+1} birds, all but one of which are non-fliers. The number of birds in the final cage that are required to have it come out that most of the time birds from the aviary are fliers is (approximately) 10^{3n} .

References:

- Bacchus, Fahiem:(1988) "A Heterogeneous Inheritance System Based on Probabilities," mimeo.
Etherington, David W.: (1987) "More on Inheritance Hierarchies with Exceptions: Default Theories and Inferential Distance," AAAI-87, Morgan Kaufman, Los Altos, pp 352-357.
Hempel, Carl G.:(1968) "Maximal Specificity and Lawlikeness in Probabilistic Explanation," *Philosophy of Science* 35, 1968, 116-133.

- Kyburg, Henry E., Jr.: (1988) "Probabilistic Inference and Non-Monotonic Inference," Shachter, Ross, and Levitt, Todd (eds): *The Fourth Workshop on Uncertainty in Artificial Intelligence*, pp 229-236.
- Kyburg, Henry E., Jr.: (1983) "The Reference Class," *Philosophy of Science* 50, pp 374-397.
- Neufeld, Eric, and Poole, David: (1988) Probabilistic Semantics and Defaults, Shachter, Ross, and Levitt, Todd (eds): *The Fourth Workshop on Uncertainty in Artificial Intelligence*, pp 275-282.
- Poole, David L.: (1985) "On the Comparison of Theories: Preferring the Most Specific Explanation," *IJCAI 85*, Morgan Kaufmann, Los Altos, 144-147.
- Reichenbach, Hans: (1949) *The Theory of Probability*, University of California Press, Berkeley and Los Angeles.
- Touretzky, D. S., Horty, J. F., and Thomason, R. H.: (1987) "A Clash of Intuitions: the Current State of Non-Monotonic Multiple Inheritance Systems," *IJCAI-1987*, 476-482.
- Touretzky, D. S.: (1984) "Implicit Ordering of Defaults in Inheritance Systems," *AAAI-84*, 322-325.

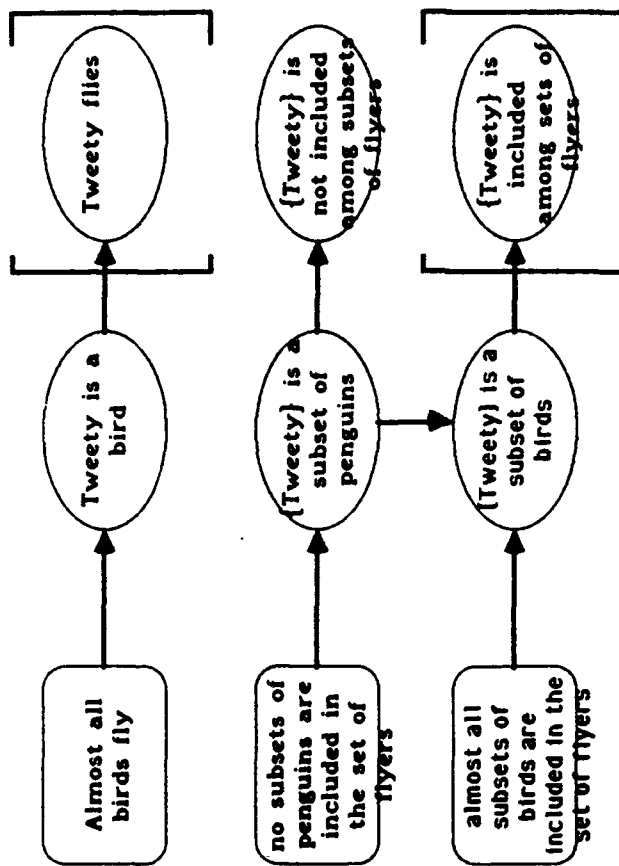
The Classical Case



The first inference is defeated by the second, since part of our knowledge includes the fact that penguins are a subset of birds.

figure 1

Birds by Other Names



The first pattern gives way to the second, not because the second refers to a subset of the birds referred to in the first (it doesn't), but because the first pattern is reflected in the third, and the third is defeated by the second according to the subset principle.

figure 2

Cages of Birds





figure 3

How Many Birds Fly?




figure 4